

Challenges with developing a Commercial P2P System

Aaron Colwell
RealNetworks

Outline

- Why P2P?
- RN Content Delivery Profiles
- Design Challenges for a Commercial P2P System
- Insights from our P2P Bandwidth Savings Study
- Future Research Directions



Why P2P?

- Economical, scalable content delivery
 - Reduction in bandwidth costs and server capacity
 - BitTorrent claims of 80+% savings are hard to ignore.
 - Availability increases with popularity -> no overprovisioning for flash crowds.
 - Large-scale events require us to maintain large server farms for events that only happen a few times a year.



Content Delivery Profiles

- Have several different delivery profiles to address
 - On-demand music service (Rhapsody)
 - Millions of clips
 - Typically 3-7 minutes, ~4-6 MB each
 - Large Live Broadcast events (Real Broadcast Network)
 - Big Brother
 - Sub-10 second latency, w/ 10s of thousands of users.
 - Movie & Casual Game Downloads (Film.com, RealArcade)
 - 10-100s of MB
 - 100-1000s of titles.



Challenges for P2P networks

- NAT Traversal
- Content Integrity
- System Security
- Churn
- Fairness
- Peer Heterogeneity
- Quality of Service
- Participation Incentives



Legal vs Illegal Content Challenges

- Illegal Content
 - Best effort service acceptable.
 - Don't expect QoS guarantees.
 - No financial investment in content.
 - Willing to take chances with potentially malicious software to gain access to the content.
- Legal Content
 - Expect a base level of QoS.
 - Content must always be available, esp. if they are paying money.
 - Customers wary of “unnecessary connections” from commercial products. Acceptance requires consumer education & incentives.
 - Content & delivery network must be secured to keep rights holders happy.



Where do we start?

- Focus on subset of the challenges that capture important aspects of our various delivery profiles.
 - Churn
 - Peer Heterogeneity
 - Quality of Service
- Determine how these challenges affect BW savings in a peer assisted CDN


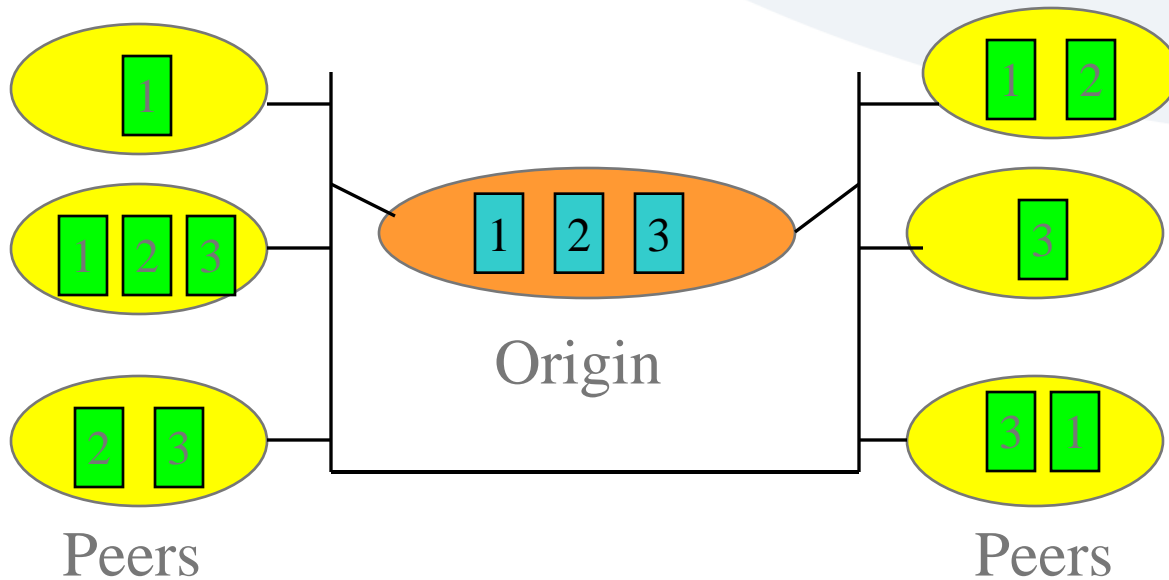


P2P Bandwidth Savings Study

- Study BitTorrent to understand potential BW savings for RN workloads.
- Explore dimensions that likely affect BW savings
 - Peer BW heterogeneity
 - Arrival/Departure processes
 - Seeding Strategies
- Determine whether BW saving are worth the effort of developing a P2P delivery system.



Evaluation Environment



Host



Link w/ varied upload/
download capacity



BitTorrent client for file n



BitTorrent tracker + seed
for file m

Multiple BT clients (peers) run across 6 hosts



WARNING

Graphs are intended to display trends and provide intuition about behavior.

The data is preliminary and should not be taken as actual BW savings.



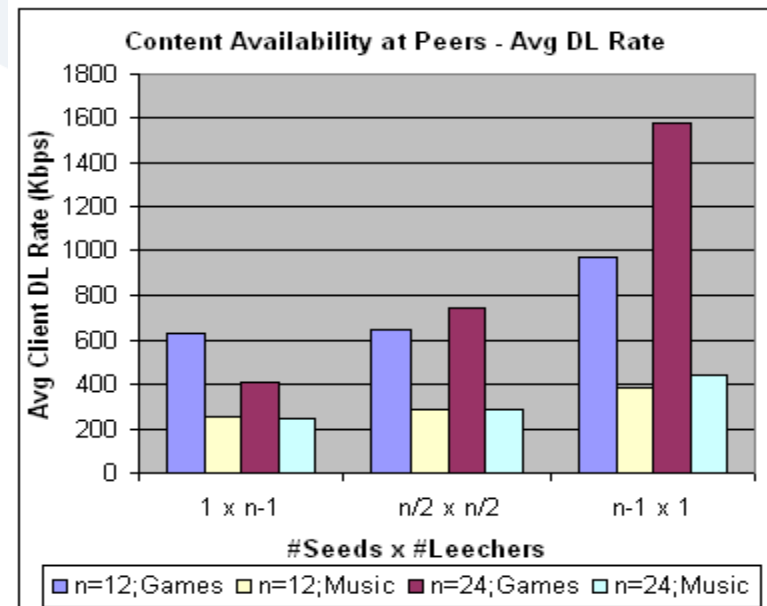
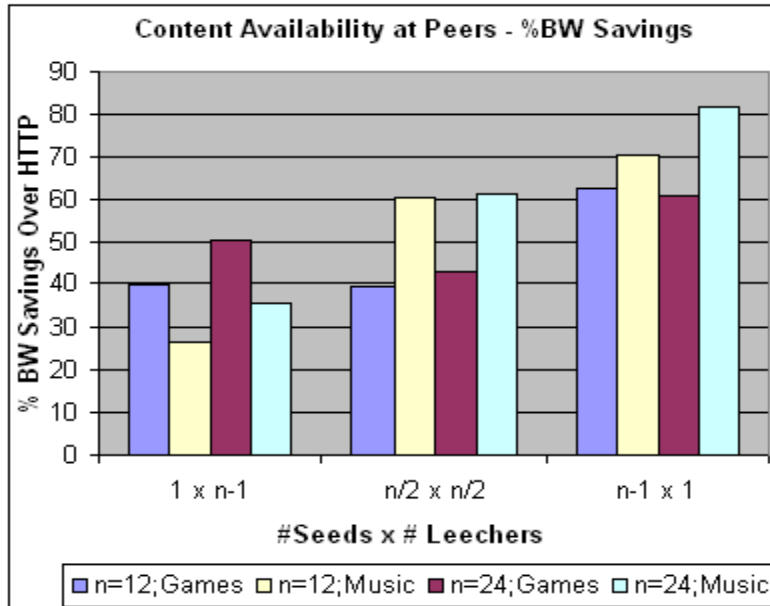
BW Savings- Impact of peer UL/DL ratio



- For fixed UL rate, as DL rate ↓, BW savings ↑
 - Lower DL rate → longer transfers & fewer peers to saturate link → more data from peers instead of origin





Impact of Content Availability at Peers*



- Content availability at peers influenced by peer seed time and file inter-reference time
 - Both factors can be captured by #Seeds:#Leechers
- BitTorrent tends to favor downloading from seeds
 - Previous results used 1 seed and n -1 leechers (i.e., 'worst' case)
- More seeds → better availability → more BW savings
 - Provide incentives to seed (inherent for live content)



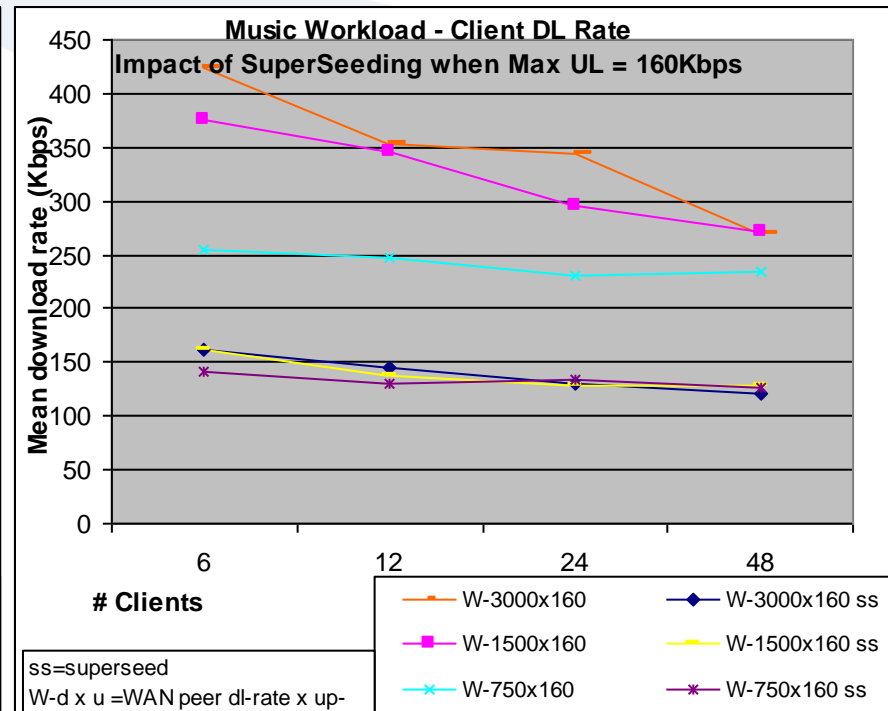
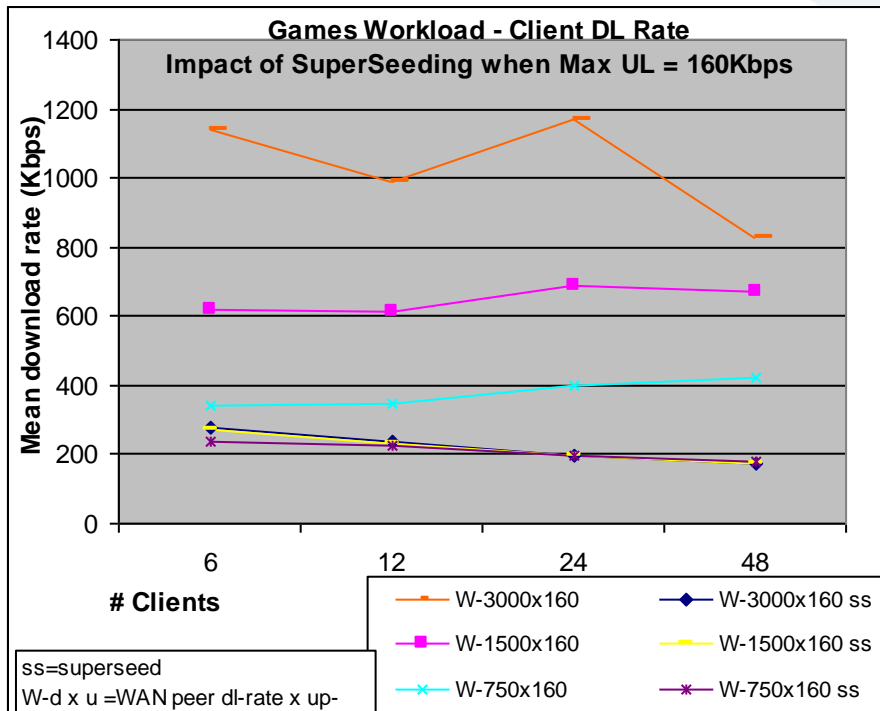
BW Savings - Impact of seeding strategies



- “Smarter” seeding to minimize BW utilization at origin
 - [SuperSeeding](#) mode (origin masquerades as leecher)
 - Explicitly cap upload rate at origin
- Significant savings in BW w/ superseeding
- But ...



Mean DL rates w/ SuperSeeding



- Mean DL rate at clients significantly lower w/ superseeding
 - Often < file encoding rate (e.g., <150 Kbps for music)
- Origin cannot attempt to reduce BW too aggressively if QoS matters



Study Conclusions

- Significant savings can be realized for a variety of workloads.
- Peer BW, mesh composition, and seeding strategy have complex interactions that vastly affect bandwidth savings.
- Key parameters need to be identified to help control BW savings.



Future Research Directions

- Further explore the parameter space covered in the study.
- Study how fairness, security, resource utilization, etc. affect BW savings.



Questions

